Paleobiology, 31(2, Supplement), 2005, pp. 146–156

The evolution of complexity without natural selection, a possible large-scale trend of the fourth kind

Daniel W. McShea

Abstract.—A simple principle predicts a tendency, or vector, toward increasing organismal complexity in the history of life: As the parts of an organism accumulate variations in evolution, they should tend to become more different from each other. In other words, the variance among the parts, or what I call the ''internal variance'' of the organism, will tend to increase spontaneously. Internal variance is complexity, I argue, albeit complexity in a purely structural sense, divorced from any notion of function. If the principle is correct, this tendency should exist in all lineages, and the resulting trend (if there is one) will be driven, or more precisely, driven by constraint (as opposed to selection). The existence of a trend is uncertain, because the internal-variance principle predicts only that the range of options offered up to selection will be increasingly complex, on average. And it is unclear whether selection will enhance this vector, act neutrally, or oppose it, perhaps negating it. The vector might also be negated if variations producing certain kinds of developmental truncations are especially common in evolution.

Constraint-driven trends—or what I call large-scale trends of the fourth kind—have been in bad odor in evolutionary studies since the Modern Synthesis. Indeed, one such trend, orthogenesis, is famous for having been discredited. In Stephen Jay Gould's last book, *The Structure of Evolutionary Thought,* he tried to rehabilitate this category (although not orthogenesis), showing how constraint-driven trends could be produced by processes well within the mainstream of contemporary evolutionary theory. The internal-variance principle contributes to Gould's project by adding another candidate trend to this category.

Daniel W. McShea. Department of Biology, Duke University, Box 90338, Durham, North Carolina 27708-0338. E-mail: dmcshea@duke.edu

Accepted: 18 August 2004

There may be no marvel greater than the complex functionality of organisms: the precise coordination of their parts, their reliability in the face of environmental upset, and their ability to act appropriately in their own interests. And then there's the sheer number of their parts and part types, across a range of physical scales, most obvious in the larger multicellular organisms. These properties together have been aptly called ''adaptive complexity'' (Ruse 2003), and the apparent increase in the adaptive complexity of organisms over the history of life is one of the great mysteries of biology.

The mystery is half solved by breaking the notion of adaptive complexity into its two components, adaptation and complexity, and thinking about them separately. Notions like ''coordination,'' ''reliability,'' and ''appropriate action'' (from the first sentence) are all aspects of function, or adaptation. And function or adaptation in organisms is no mystery at all. Since Darwin, we have had a perfectly adequate explanation for it, natural selection.

Complexity is another matter. Suppose we do a thought experiment in which we subtract all aspects of adaptation from our understanding of an organism. In other words, we put aside everything we know about how the organism functions. What's left is an assemblage of parts, diverse in their sizes, compositions, shapes, and orientations with respect to each other. To capture that diversity, we will call the organism complex, although keep in mind that this is what you might call ''pure complexity,'' divorced from any notion of function. For complexity in this sense, biology has no satisfactory explanation. Consider just complexity in the sense of number of part types. Why should organisms have so many? And why would complexity in this sense increase in evolution (as the common view suggests it does)? In other words, why would modern organisms be more complex than ancient ones?

One possibility is that natural selection has favored complexity along with functionality. Perhaps functional improvement requires

 0094-8373/05/3102-0011/$1.00

greater division of labor, which in turn requires more part types, in other words, greater complexity. Or maybe evolutionary increases in body size—also favored by selection perhaps—demand greater complexity for functional reasons (Bonner 1988). (For example, very small organisms can breathe by diffusion but larger ones need to add a circulatory system.) These arguments are plausible, and others equally plausible can be imagined. But they may not be necessary.

Here I argue that organisms are expected to accumulate variations spontaneously as they evolve, with the result that their internal parts become more differentiated. In other words, there is a spontaneous tendency for what I call their ''internal variance'' to increase. I further argue that internal variance is an aspect of complexity, and therefore the internal-variance principle generates a vector in evolution toward increasing complexity. The internal-variance principle is quite general. It explains, for example, the increase in differentiation of a newly painted picket fence as the paint on each picket acquires its own unique pattern of wear, and the increase in differentiation of the surface of the moon as it accumulates impact craters.

The vector is a generative tendency, or a bias in the production of variants, and therefore it predicts a trend, at least prior to any consideration of selection. But I call it a vector rather than a trend, because the evolutionary resultant is unknown. Selection could reinforce the internal-variance vector, act neutrally, or oppose it. If it acts in opposition, selection might even overwhelm the vector, either precisely canceling it or producing a net vector toward decrease.

In the following discussion, I do not challenge the conventional wisdom that a long-term trend in mean complexity—understood as differentiation of parts—in fact occurred, at least in metazoans (although later I will neutrally consider what it would mean for the internal-variance principle if we were to discover that no trend in the mean had occurred). It is worth noting in this context that some evidence for a metazoan trend exists (e.g., Valentine et al. 1994) but that doubts have been raised, especially for a trend after the Cambrian explosion (e.g., Simpson 1967; McShea 1996).

In a classification of causes of evolutionary trends, the internal-variance principle falls in with some heretical company. As will be seen, if the principle is correct, the vector it predicts should be pervasive, affecting all lineages in all times and all places in the history of life, and therefore would constitute a kind of evolutionary ''drive'' (sensu McShea 1994; Gould 2002). And any trend resulting must be considered driven, not by selection but by a kind of mathematical or statistical constraint.

Constraint-driven trends—when they occur in higher taxa over substantial time spans—fall into a category that I call large-scale trends of the fourth kind (explained later in a discussion of a four-part trend-classification scheme). This category includes some, such as orthogenesis, that are notorious in the history of biology for having been discredited. In his last book and magnum opus, Stephen Jay Gould (2002) tried to rehabilitate constraint-driven trends (although not orthogenesis itself), offering a number of real and hypothetical examples and showing that they could be produced by causes falling well within the present theoretical mainstream. The internal-variance principle contributes to this rehabilitation, adding one—and I believe an important one, at the highest level of generality—to Gould's list of examples.

The internal-variance principle has a long and diverse ancestry. It comes mainly from Herbert Spencer's (1900, 1904) metaphysics, especially his notion of ''the instability of the homogeneous'' (see discussion in McShea 1991), which he invoked to explain organismal complexity specifically. Thus my treatment here can be understood as an attempt to revive, update, and further explore Spencer's notion. The principle can also be found in other guises. It is closely related to another old notion in biology, the duplication and differentiation of parts (e.g., Gregory 1935; Weiss 1990). And it is implicit in modern treatments of morphological evolution as a diffusive or Markov process (e.g., Raup 1977). Finally, what is essentially the same principle has been recognized in a related context, the evolution of differences among species (rather than as

here, differences among parts in an organism). In particular, it underlies the intuition that the degree of morphological differentiation among species in a group, or their ''disparity'' (Foote 1997), should tend to increase spontaneously (Ciampaglio et al. 2001). And it emerges in the writings of one of the thermodynamic schools of thought in evolutionary studies (e.g., Wicken 1987; Brooks and Wiley 1988), where it has been invoked mainly to explain species diversity.

The Internal-Variance Principle

Figure 1 illustrates the principle with a simple model. Five parts of an organism are plotted on an axis representing some dimension, say, their length (actually log-length). Initially all five parts are identical (e.g., serial homologs, perhaps), with their lengths set at the same value, arbitrarily, 10 millimeters. In the first time step, random heritable variation is added to each part, so that its length increases or decreases by the same factor, divided by 0.9 for increases and multiplied by 0.9 for decreases, with the direction of change chosen at random (probability of increase = probability of decrease = 0.5). Each part is treated independently. And in each subsequent time step, more variation is introduced, and new values calculated for each of the five parts based on their values in the previous time step, so that change is cumulative.

The top box in Figure 1 shows the starting distribution of logged part-lengths in a single run of the model. The parts have the same lengths so the points fall precisely on top of each other. The lower boxes show the distributions after 1, 2, 5, 10, and 20 time steps. The figure also shows the increase in internal variance quantitatively. The term ''variance'' in this phrase is intended in a general sense, to mean something like ''amount of variation'' or ''degree of differentiation,'' not in a formal statistical sense to refer to a sum of squared deviations from a mean. However, statistical variance is one possible measure of internal variance. In Figure 1, internal variance is measured as the square root of the statistical variance, i.e., the standard deviation.

Figure 2A shows the trajectory of the standard deviation over the entire run, and Figure

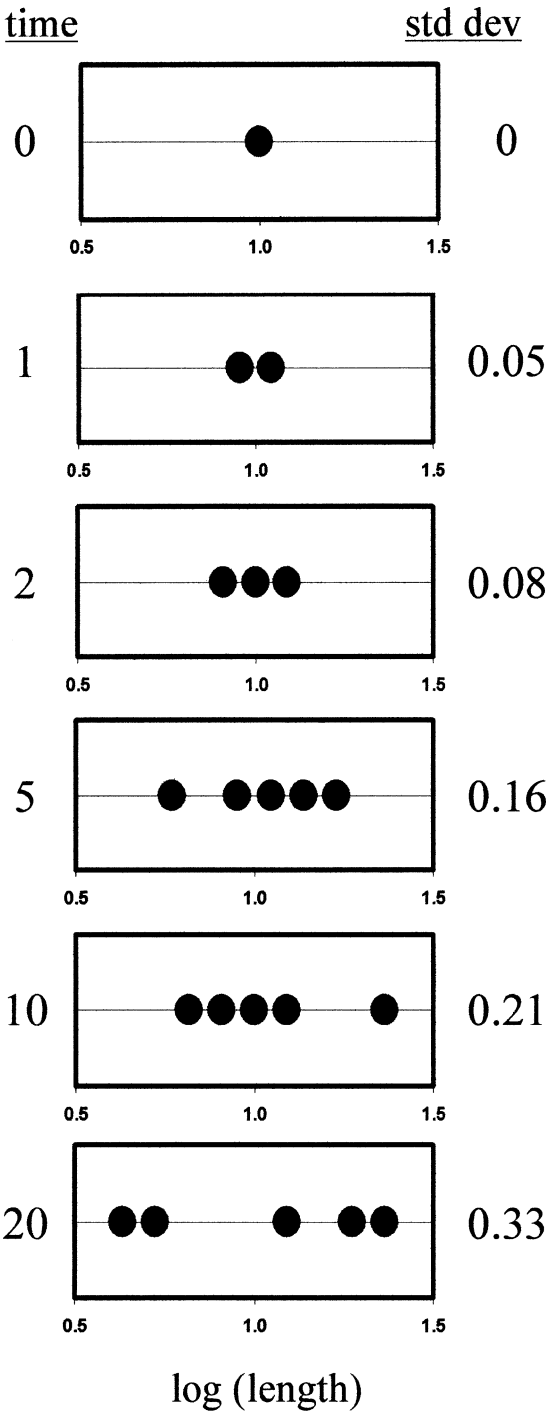

FIGURE 1. The internal-variance principle. The dots are five initially identical parts (overlapping in the top frame). In each time step, a variation introduced to each part changes its length by a factor of 0.9 (length is multiplied by 0.9 for decreases and divided by 0.9 for increases), with increases and decreases occurring with equal probability (0.5). Internal variance is measured as the standard deviation among log-lengths. The numbers on the right show the internal variance at time 0 (top frame), 1, 2, 5, 10, and 20 (bottom frame).

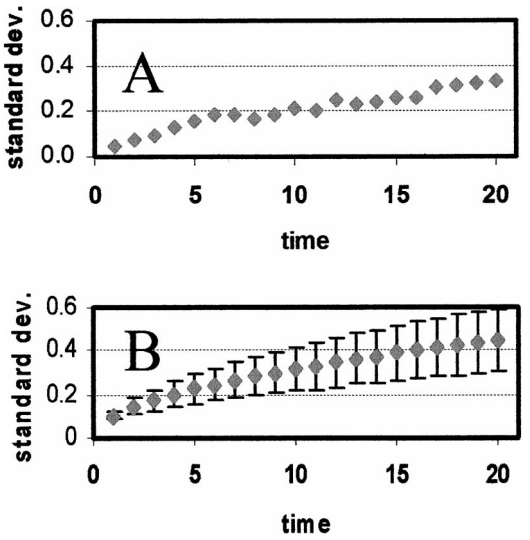

FIGURE 2. A, The trajectory of the standard deviation for a single run of the model in Figure 1 over 20 time intervals. B, The trajectory averaged over 1000 runs. Error bars show one standard deviation of the internal variance (i.e., the standard deviation of the standard deviation).

2B shows the mean standard deviation over 1000 runs, with error bars (i.e., the standard deviation of the standard deviation). Analytically, it can be shown that the internal variance rises as the square root of the number of time steps. Notice that the internal variance rises even after parts have become highly differentiated. That is, even when complexity is high, the expectation is further increase, without limit.

The explanation for the trend should be obvious. In a very small number of cases, the introduced random variations may by chance make the parts more similar to each other, but in the vast majority of cases, they will cause the parts to diverge, to differentiate. In other words, there are many more ways for internal variance to increase than for it to decrease, and therefore most combinations of random variations will produce increases.

This example is a paradigmatic case. But the principle should be far more general. In particular, it should apply in the following cases: (1) To any organism with some degree of modular construction, i.e., any organism with parts (although an analogous principle could be developed for a hypothetical organism without any clear compartmentalization). (2) To any set of parts within an organism sharing a common dimension, not just initially identical parts as in the example above. Thus, the variance should increase even when the parts in the set are disparate and seemingly incomparable, such as the set: foot, kidney, ear, and aorta (perhaps in a mammal). All of these share a dimension, length for example, and therefore the variance in that dimension should increase over time. (3) To any dimension or set of dimensions. For those four mammalian parts, the variances among their lengths, widths, permeabilities, color, metabolic rates, etc. should all increase over time. And if internal variance is measured in a multidimensional space, the overall rate of increase should be higher with greater dimensionality. (4) For all intuitively reasonable measures of internal variance. (5) To cases where discrete part types are recognized. To see this, imagine that the simulation above were run with the length axis in Figure 1 divided into equal-sized bins, with each bin corresponding to a discrete part type. (6) To individuals in all lineages, over the history of life.

(7) For most types of random variation. In the example above, variation was introduced to each part independently, using a simple two-state function. However, the vast majority of arbitrary functions that could be introduced to represent variation should also produce increases in internal variance, on average, even the very complex or highly structured ones. In two other papers (McShea 1992, in press), I experimented with real and hypothetical homologous series (representing vertebral columns), to study the effect of introducing highly structured variations in which changes among adjacent parts are correlated (e.g., a sine curve), such as might be caused by introduced fields or gradients. The thinking was that such variations were at least somewhat realistic and might represent counterexamples, that is, they might tend to produce decreases in internal variance. In fact, some functions could produce decreases, but only occasionally, with certain parameter settings. Increases still predominated.

Notice that the internal-variance principle is different from another trend mechanism, what has been called variously a passive mechanism, an increase in variance, or diffusion in the pres-

ence of a lower bound (Stanley 1973; Fisher 1986; McShea 1994; Gould 1996). The internal-variance principle does invoke boundaries in the sense that variances are bounded at zero. And it invokes diffusion in the sense that the morphologies of parts within an organism diffuse in morphospace. But at the lineage level, the principle predicts a pervasive increasing tendency. It predicts that parts should vary so that differences among them increase in essentially all lineages, from the simplest to the most complex. In other words, the trend in complexity should be driven (McShea 1994; Gould 2002). In contrast, a passive mechanism predicts equal numbers of lineages increasing and decreasing.

Internal Variance Is a Type of Complexity

I cannot give the full argument here for this claim (but see McShea 1996, in press). But briefly, the argument is based on a view of complexity as the amount of differentiation among parts, where variation is continuous, or as number of part types, where variation is discrete.

Some objections will arise. First, in this view, complexity is purely about structure, but colloquially the word includes some notion of function as well. We commonly think of complex systems not just as differentiated, but also capable. Now this compound usage may be useful in certain contexts, but it confounds the two central issues involved in adaptive complexity: adaptation (i.e., function), for which we have an almost universally acknowledged explanation, and structural complexity, for which we don't. And to address them separately, it is essential to use conceptually nonoverlapping terms. The case for conceptual separation is bolstered by the ease with which we can imagine objects with any of the four combinations of properties: (1) functional and structurally simple (e.g., virus, screwdriver); (2) functional and structurally complex (*Amoeba*, vacuum cleaner); (3) nonfunctional and simple (grain of sand, drop of water); and (4) non-functional and complex (dead *Amoeba*, surface of the moon).

A second objection is that even when used in a purely structural sense, complexity sometimes refers to more than just internal variance. There is also complexity in the sense of degree of hierarchical structure (number of levels of parts within wholes), irregularity of spatial configuration, and so on. The argument here is that the internal-variance principle predicts a rise in complexity only in the sense of differentiation among parts (or number of part types), what I elsewhere call ''nonhierarchical object complexity'' (McShea 1996). However, the argument is extendable to make analogous predictions at least for irregularity of spatial configuration, and perhaps for other senses of complexity as well. (At present, though, I do not see how it could predict hierarchy.)

Third, this view might seem to equate complexity with entropy, and indeed entropy is the term some have chosen for this same principle in other contexts (e.g., Brooks and Wiley 1988). Entropy in its information-theoretical interpretation refers to a relationship between microstates and macrostate, namely the more entropic macrostate is the one with the greater number of microstates corresponding to it. This is what makes the more entropic macrostate the more probable one and the one toward which the system spontaneously evolves. My house is more likely to become messier than neater, because many more configurations of my stuff (microstates) correspond to messier (a macrostate) and relatively few correspond to neater (a macrostate). Similarly, my argument here is that many more character state distributions correspond to greater internal variance than correspond to less, so greater internal variance is more probable, and it is toward greater internal variance that the organism spontaneously evolves (considerations of selection aside). So entropy would seem to be the right word.

Confusion arises however, because the word in common usage also connotes disorganization and loss of function. We think of systems as falling apart, and failing to work, as they become more entropic. But this connotation runs at a right angle to the technical meaning, which is completely silent on such matters. My messy house may be quite functional, for my purposes, even if it has high entropy. Perhaps various messy distributions of items allow me to find and use them efficient-

ly. Likewise, an organism may have high entropy, with many disparate parts, irregularly configured, and yet it may function quite well. That is, it may be quite fit. Technically, nothing in the notion of entropy conflicts with the notion of functionality, or even with increase in functionality. But given the friction between technical meaning and common-usage connotation, I have opted against entropy in favor of a more neutral term, internal variance.

To conclude this section, I understand internal variance as a type or subcategory of complexity. However, in what follows, I will risk causing some confusion by using the terms interchangeably (in effect ignoring the many other senses of complexity). The point is to subtly reinforce the connection between two notions that might formerly have seemed unconnected—one a variable that is unproblematic and even measurable, internal variance, and the other a concept that has historically been very troublesome, complexity.

The Effect of Selection

The internal-variance principle describes a tendency for the accumulation of variation to produce greater complexity, prior to any consideration of selection. When we do take selection into account, there are four ways in which it could act: (1) A variation producing a complexity increase could have a greater probability of being favored than one producing a decrease, and in that case selection will just reinforce the internal-variance vector. (2) Selection could be neutral with respect to complexity. That is, suppose that every introduced variation is favored or rejected depending on its effect on adaptedness, but that there is no correlation between adaptedness and complexity. In that case, the few variations that pass the selective filter will be a representative sample of those generated, and therefore will be more complex, on average. It is as if selection were picking from a deck that has been stacked with black cards, where the black cards represent high complexity and the red represent low. If selection has no color preference (i.e., if it is neutral with respect to complexity) and if the aces are the well-adapted forms, then most of the aces selected will be black. In other words, most of the well-adapted forms will be complex.

(3) Selection could oppose (but not overwhelm) the internal-variance vector, meaning that it could disproportionately favor the rare variants producing decreases in complexity (or disproportionately oppose the many variants producing increases). There are at least two ways it might do this. First, selection might impose boundaries on variation, so that increase in the length, permeability, etc. of parts beyond some upper limit, and decrease below some lower limit, is disfavored. Such boundaries must exist in that all variation has functional limits. But notice that for boundaries to limit internal variance, the distribution of parts would have to be already maximal, filling the entire functionally permissible range. Also, maximal dispersion must have already been reached in all dimensions. If dispersion can continue in even one dimension, then overall internal variance should still increase.

Alternatively, it could be that simpler species have disproportionately more ecological opportunities, on average, than complex species. In other words, the generative bias produces more increases in internal variance than decreases, but this vector is partly offset by the disproportionate frequency of ecological opportunities, of niches or places in the economy of nature, for simpler species. This might seem implausible if ecological opportunities for increased complexity spring to mind more readily than those for decrease. For example, specialization is probably a common route to adaptation, and specialization might require greater complexity, on average. Arguably, however, opportunities for simplicity should also be common. Obvious candidates include the opportunities associated with the evolution of parasitism or of small size generally (as in the evolution of the so-called interstitial fauna or meiofauna). But the range is really much broader. For example, consider the evolution of the mammalian vertebral column in the transitions from terrestrial to aquatic living. Terrestrial ancestors typically had complex columns suitable for quadrupedal locomotion on land (e.g., with special modifications of the column for attachment of the hind

legs). But their aquatic descendants, such as whales, tended to have simpler, more uniform, fishlike columns, suitable for undulatory propulsion in the water (McShea 1991). Selection favored simplicity, not for any reason having to do with simplicity itself, but on account of an accidental association between simplicity and a particular functional mode. Such associations must be quite common.

(4) Finally, selection might oppose the internal-variance vector—by either of the routes in case 3 above—strongly enough to cancel it or even to produce a reverse vector. Notice that for cancellation, selection must allow increases and decreases in equal numbers, not equal proportions (which would be equivalent to selective neutrality).

Cases 1, 2, and 3 predict a driven trend in mean complexity, produced either partly (case 1) or entirely (cases 2 and 3) by the internal-variance principle. Case 4 predicts no trend in the mean if selection precisely cancels internal variance, at least no driven trend (the mean could still increase by a passive mechanism). And it predicts a decreasing trend if selection is strong enough to produce reversal.

Unfortunately, we do not know the fact of the matter, for either the mean or the mechanism, for life as a whole or for any kingdom-level clade within it. Several cases have been studied within the metazoans, and both the passive and driven mechanisms have been found (McShea 1993; Saunders et al. 1999; Sidor 2001). The best that can be done at this point is to frame a conclusion in the conditional. In a given group, at some specified temporal scale, if no trend in the mean occurred, or if a trend occurred and is ultimately discovered to have been passive, then presumably selection opposes and cancels the internal-variance vector. But if mean complexity increases, and if the trend is ultimately found to be largely driven, then the internal-variance principle must account for it, at least partly. It would still be impossible to say whether the trend was occurring with the support, neutrality, or partial opposition of selection (cases 1, 2, and 3, above). To infer that, we would need to know the magnitude of the internal-variance vector.

One other possible selective route needs to be considered, selection above the level of the organism, such as species selection. For example, it might be that species with complex morphologies are more specialized and therefore more extinction prone. Again, however, nothing conclusive can be said. Schopf et al. (1975) investigated this in metazoans and found a positive correlation between complexity and extinction probability but interpreted it as an observational artifact.

A final point about case 4: from one viewpoint, the impact of cancellation on our understanding of the passive mechanism is fairly staggering. Prior to any consideration of internal variance, it is easy and tempting to see the passive mechanism as a kind of null case, one in which no forces act at the large scale. (That's why increases and decreases are equally frequent.) But the internal-variance principle forces us to reconceptualize, to see a passive trend as necessarily the resultant of two equal and opposite vectors, internal variance and selection, both potentially quite strong. Indeed, consider the enormous force that selection must exert in order to cancel the internal-variance vector. If the vast majority of variations arising are increases in complexity, then selection must oppose virtually all of them, while favoring a comparatively large proportion of the small number of decreases arising, in order to leave equal numbers of each. To put it another way, ecological opportunities for more-complex variants must be vanishingly rare, or improbable, compared with those for less-complex variants. Thus, if there was a trend in mean complexity over the history of life, and if we discover that it was passive, we are left with a real puzzle: Why does selection abhor complexity?

Objections

I foresee three objections to the internal-variance principle. First, there is a common notion that losses of parts—and therefore losses of complexity—are more likely than gains because most mutations affecting a given developmental pathway will tend to disrupt it. I think this intuition could be partly right and, insofar, could be fatal to the internal-variance principle. It is right in that truncations of development are probably developmentally easy

and can produce reductions in internal variance. Internal variance certainly increases in ontogeny, because organisms go from simple to complex as they develop, and therefore truncations should, on average, leave structure simplified overall. But intuitively anyway (and in the absence of any organized data on the subject), this should be the case only for major truncations that cut short all of development or remove all of some quasi-independent module. Partial truncations that remove arbitrary subsets of features should not affect internal variance, on average, for the simple reason that the expected variance of a set of data (measurements of a set of structures before truncation) is the same as the expected variance of a random subset of those data (those remaining after truncation). Thus, the important empirical question for the internal variance principle is the frequency with which major or modular truncations actually occur, i.e., whether they occur often enough to offset the internal-variance vector.

The idea that loss of structure is developmentally easy may also be partly faulty. The fault lies in a common confusion between loss of function and loss of structure. Random modification must tend toward loss of function, but need not tend toward loss of structure. It is true that a random variation might lead to structural loss via a failure of an inductive tissue contact, elimination of a differentiation event, or a failure of the production of a tissue fold. But given the dynamism of development, would a random variation not be equally likely to lead to structural gain, perhaps by producing a loss of an inductive barrier, an extraneous differentiation, or the failure of a fold-smoothing-mechanism? Are not developmental aberrations, ''monstrosities,'' often more complex structurally? I think these are currently open questions.

Another objection, one might argue, is that selection has favored mechanisms that buffer or canalize development against most of the variations that would increase complexity, because these are mostly deleterious. Thus, arguably, most increases in complexity will never be expressed, perhaps negating the internal-variance vector. However, decreases in complexity are also likely to be deleterious. What buffering mechanisms oppose is departure from a functional norm, not greater complexity per se, and therefore they should oppose the expression of both increases and decreases equally.

Finally, one might object that direct testing of the internal-variance principle is difficult, perhaps impossible. The principle predicts that offspring should be more internally varied than their parents, but only prior to any consideration of selection. And it may be impossible to find circumstances in which selection is truly absent. However, we might look at changes in internal variance in special cases, where the effect of selection should be minimal. For example, we might measure some dimension of each of the teeth in a mammalian adult tooth row and compare the standard deviations of this measure in a large sample of parent-offspring pairs. The teeth do not emerge and function until adulthood, and so arguably, before they do, they ought to be less subject to selection than other structures. Comparable measurements could be obtained from parents and offspring at the same life stage, ideally when the teeth are just erupting. The internal-variance principle predicts that the standard deviation for the offspring will be higher, on average. Of course, the test is imperfect in that teeth are undoubtedly connected to other structures in development and are presumably subject to indirect selection by that route.

Large-Scale Trends of the Fourth Kind

Figure 3 shows a classification of large-scale trends based on two dichotomies, one between alternative mechanisms (rows), and one between alternative underlying causes (columns). A large-scale trend is a long-term directional change in a summary statistic for a clade, such as the mean. Mechanism refers to the pattern of change in the variable in question, that is, whether the trend occurs because of a pervasive tendency in one direction or the other (i.e., a driven trend) or a local inhomogeneity (i.e., a passive trend), such as a boundary (McShea 1994). Causes occur at a lower level, answering the question of whether the drive or boundary is the result of selection or of any of the various types of constraint (e.g.,

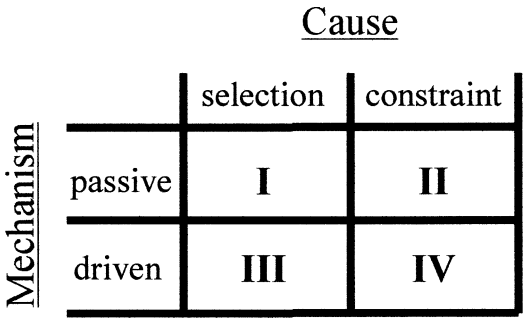

FIGURE 3. A classification of large-scale trends, according to both mechanism (rows) and underlying cause (columns). See text.

Gould 1989, 2002). (This form of the passive-driven distinction is a crude one; see Alroy 2000; McShea 2000; and Wagner 1996 for more-sophisticated alternatives.)

The four quadrants represent the possible combinations of mechanism and cause, and indeed it seems likely that all are occupied by actual evolutionary trends at some scale. Here I will consider only trends in complexity. Trends in the first quadrant, or ''large-scale trends of the first kind,'' are passive with selection as the underlying cause. These include trends in which increases and decreases occur equally frequently and there is a selective lower boundary. For complexity, such a boundary might arise if organisms less complex than some minimum are not viable. It might be, for example, that life processes cannot be carried out with fewer part types than in a prokaryotic cell, and therefore variants arising at a lower complexity, below the boundary, are selected against. The resulting increase in mean complexity over the history of life would be passive, with selection as the underlying cause.

Large-scale trends of the second kind are passive with constraint as underlying cause. For example, the number of cell types in a eukaryote cannot be less than one, and this mathematical (or logical) constraint imposes a lower bound on number of cell types in metazoans. Valentine et al. (1994) suggested that the increase in number of cell types over the history of animals may have been passive in this way.

Large-scale trends of the third kind are driven with selection as the driving force. The increase in septal-suture complexity in Paleozoic ammonoids may have been a trend of this sort. Saunders et al. (1999) documented that suture complexity increased significantly more often than it decreased among lineages, and they speculated that selection may have been the cause of this upward bias.

Large-scale trends of the fourth kind are driven, so that increases occur more often than decreases among lineages, with constraint producing the upward tendency. Such trends were not well received in evolutionary studies over most of the twentieth century, largely on account of their historical association with orthogenesis. The more extreme of the nineteenth-century orthogeneticists argued that evolution proceeds by acceleration of ontogeny, leaving room for terminal addition of a predictable sequence of increasingly senescent developmental stages. The principal objection to notions of this sort, from the perspective of the Modern Synthesis, is that they seem to justify interpreting certain evolutionary results—such as the massive antlers of the Irish elk—as maladaptive (Gould 2002). But Gould argued that there is nothing in the nature of constraint-driven trends that forces such interpretations, especially if the constraint occurs at the species level. He offered a possible example (citing Strathmann): species with non-planktotrophic larvae can arise from planktotrophic species by the loss of larval swimming ability, but the reverse is much less likely, because regaining a swimming larva is developmentally more difficult, i.e., constrained. The result is not maladaptive design in individuals, just a bias toward the production of non-planktotropic species.

The internal-variance principle is a consequence of mathematical or statistical necessity and therefore constitutes a kind of constraint (what Gould [1989] called a ''formal'' constraint). And the constraint is expected to operate in all lineages, so any resulting trend must be driven. A difference with the planktotrophy example is that the internal-variance vector acts at the organism level, as in orthogenesis, not the species level. But unlike orthogenesis, the driven variable is a higher-level or abstract property of an organism, its internal variance. So the principle dictates no

specific change in structure. And because all variants must pass through the selective filter, the specific structures that survive should all be adaptive.

Conclusion

The internal-variance principle identifies an evolutionary vector, a kind of pervasive force pushing complexity—understood as differentiation among parts—upward. But it is a separate matter whether this vector is actually manifest as a bias in the direction of change in evolutionary lineages, and whether there is an actual trend in, say, the mean complexity of metazoans. If there is a trend, and if it is driven, then the internal-variance vector supplies a likely cause (with or without reinforcement by selection). Or the vector could be opposed by selection to the point where it is overcome, with the result that either no trend or only a passive trend occurs.

Regardless of whether the internal-variance vector is manifest, its existence demands what for some will be a dizzying reversal of intuition. We are used to thinking of complexity as hard to produce. The internal-variance principle shows it to be easy, in particular to be easy to generate in development and therefore in evolution. To see this, we have to purge the conventional notion of complexity of its functional component, and to conceive complexity as purely structural. That is, we have to see complexity—in the relevant sense here—as just structural diversity or differentiation among parts, i.e., internal variance. Having done so, it is easy to see why it arises spontaneously. And it becomes equally easy to see why simplicity, complexity's opposite, is developmentally hard (except perhaps via major truncations of development).

Importantly, the point is not that *adaptive* complexity is easy, not that specific complex structures like eyes and brains spontaneously self-assemble somehow in evolution. Rather, it is that most spontaneous changes produce greater complexity. Of these, the vast majority will be nonfunctional. Eyes and brains represent the tiny subset that are functional. Adaptation stills needs selection. But complexity does not.

Acknowledgments

The idea in this paper owes much to the exhilarating metaphysical thinking of H. Spencer, but also to the more recent inspirational ideas of J. T. Bonner, D. Raup, L. Van Valen, and my colleagues in the biology and philosophy discussion group here at Duke University. Many thanks also to A. Hallam and an anonymous reviewer for useful critiques. I regret that I did not get a chance to run this idea by Stephen Jay Gould. But I want to thank him anyway, for his exuberant iconoclasm, and for writing so well and thinking so clearly about evolutionary trends over the years.

Literature Cited

Alroy, J. 2000. Understanding the dynamics of trends within evolving lineages. Paleobiology 26:319–329.

Bonner, J. T. 1988. The evolution of complexity. Princeton University Press, Princeton, N.J.

Brooks, D. R., and E. O. Wiley. 1988. Evolution as entropy: toward a unified theory of biology. University of Chicago Press, Chicago.

Ciampaglio, C. N., M. Kemp, and D. W. McShea. 2001. Detecting changes in morphospace occupation patterns in the fossil record: characterization and analysis of measures of disparity. Paleobiology 27:695–715.

Fisher, D. C. 1986. Progress in organismal design. Pp. 99–117 *in* D. M. Raup and D. Jablonski, eds. Patterns and processes in the history of life. Springer, Berlin.

Foote, M. 1997. The evolution of morphological diversity. Annual Review of Ecology and Systematics 28:129–152.

Gould, S. J. 1989. A developmental constraint in Cerion, with comments on the definition and interpretation of constraint in evolution. Evolution 43:516–539.

———. 1996. Full house: the spread of excellence from Plato to Darwin. Harmony Books, New York.

———. 2002. The structure of evolutionary theory. Harvard University Press, Cambridge.

Gregory, W. K. 1935. Reduplication in evolution. Quarterly Review of Biology 10:272–290.

McShea, D. W. 1991. Complexity and evolution: what everybody knows. Biology and Philosophy 6:303–324.

———. 1992. A metric for the study of evolutionary trends in the complexity of serial structures. Biological Journal of the Linnean Society 45:39–55.

———. 1993. Evolutionary change in the morphological complexity of the mammalian vertebral column. Evolution 47: 730–740.

———. 1994. Mechanisms of large-scale trends. Evolution 48: 1747–1763.

———. 1996. Metazoan complexity and evolution: is there a trend? Evolution 50:477–492.

———. 2000. Trends, tools, and terminology. Paleobiology 26: 330–333.

———. In press. A universal generative tendency toward increased organismal complexity. *In* B. Hallgrimsson and B. Hall, eds. Variation. Academic Press.

Raup, D. M. 1977. Stochastic models in evolutionary paleobiology. Pp. 59–78 *in* A. Hallam, ed. Patterns of evolution as illustrated by the fossil record. Elsevier, Amsterdam.

Ruse, M. 2003. Darwin and design: does evolution have a purpose? Harvard University Press.
Saunders, W. B., D. M. Work, and S. V. Nikolaeva. 1999. Evolution of complexity in Paleozoic ammonoid sutures. Science 286:760–763.
Schopf, T. J. M., D. M. Raup, S. J. Gould, and D. S. Simberloff. 1975. Genomic versus morphologic rates of evolution: influence of morphologic complexity. Paleobiology 1:63–70.
Sidor, C. A. 2001. Simplification as a trend in synapsid cranial evolution. Evolution 55:1419–1442.
Simpson, G. G. 1967. The meaning of evolution. Yale University Press, New Haven, Conn.
Spencer, H. 1900. The principles of biology. Appleton, New York.
———. 1904. First principles. J. A. Hill, New York.
Stanley, S. M. 1973. An explanation for Cope's rule. Evolution 27: 1–26.
Valentine, J. W., A. G. Collins, and C. P. Meyer. 1994. Morphological complexity increase in metazoans. Paleobiology 20: 131–142.
Wagner, P. J. 1996. Contrasting the underlying pattern of active trends in morphologic evolution. Evolution 50:990–1007.
Weiss, K. M. 1990. Duplication with variation: metameric logic in evolution from genes to morphology. Yearbook of Physical Anthropology 33:1–23.
Wicken, J. S. 1987. Evolution, thermodynamics, and information. Oxford University Press, New York.